I want to turn to another topic here, a closely related topic. And that is adaptive management. Adaptive management, there's a number of different views about adaptive management. I want to talk about a view of adaptive management that's now referred to as the decision theoretic school of adaptive management.

And the idea here is that we can reduce uncertainty while we're making decisions in some settings. So there's a special class of decisions we haven't really talked about in this course yet. And those are linked decisions that are recurrent, that occur repeatedly over time. And most commonly these are annual decisions really where there's a similar decision that we make each year.

And the interesting thing about these recurrent linked decisions is that they afford us some interesting opportunities. There's connections. There's connections between what we do this year and what we're able to do next year and the subsequent year and so on.

And these decisions give rise to some interesting opportunities, particularly when they're made in the face of uncertainty. Because we have the opportunity, perhaps, to reduce uncertainty early on and apply that new information to subsequent decisions in the future.

And so this question of how we simultaneously manage a system in the face of uncertainty while thinking about how we would reduce that uncertainty in applying it to subsequent decisions, that's an interesting problem. And it's actually a very, very common problem in natural resource management.

A lot of our decisions are recurrent in this manner. And almost all our decisions are made in the face of uncertainty. So we often have this opportunity to reduce uncertainty while making decisions.

So let's talk a little bit about that. What makes these decisions different? What are sort of the unique aspects of these decisions? And how do we need to think about them differently?

Well, these decisions have two potential characteristics. One is they're dynamic. And two is that they might be adaptive. And let me talk about these two characteristics.

What are dynamic decisions? And what makes them different? Well, dynamic decisions are ones that are linked in time like this. And there's two issues that come up. One is that today's decision affects

future decisions. It affects what happens immediately, what result we get immediately.

But it also affects where the system is going to be, the state of the habitat or the population next year when we make our next decision. So today's decision affects what's going to happen in the future and, therefore, future decisions. That's the dynamic aspect of it.

But the consequences that we get, the observation that we make subsequently, can help us reduce uncertainty. And so we have the opportunity to be adaptive as well. All right, so the dynamic part of this is that the system's in some state today, right? It's a population size or habitat conditions or weather conditions, whatever. The system is as it is today.

I take some action on it today. And I get some immediate return. But that also puts the system in a new state next year let's say, at the next time period. And that new state that it's in, it can be affected by the action I took currently. So I affect both the current return and the system state in the next time period.

I take an action next year. That generates a return next year and then puts the system in a different state two years from now and so on. And so the point is that we've got this linkage through time that the actions we take generate current returns.

But they also affect future returns by changing the state of the system next year and, therefore, changing what we might do, what action we might take. So that's the dynamic nature of recurrent decisions about how the decisions we make now affect decisions we can make and returns that we can get in the future.

These decisions also have the opportunity to be adaptive. And if you look at the diagram on page J-5 in your notebook, let's say right now our predictions about how the system works are uncertain. So this is sort of shown with this wide normal distribution that we aren't quite positive about what's going to happen if we take a particular action.

But let's say we take a particular action. We've made some predictions with our system model. We actually take an action at time now, current time, time T, let's say.

But then we also have the monitoring system, which allows us to observe the system at the next time period after our action's taken place. So we've got a prediction about what our action would have resulted in. And then we've got an observation about what it did in fact result in. And the tension

between these two things, the tension between the prediction and the observation, allows us to generate learning.

We learned something about how good a job our model was doing in making that prediction. That learning sort of, in a sense, tightens up our predictions now. We learned something about the system model. We can now make tighter predictions.

And so that's depicted in this diagram by a tighter normal distribution. We've got a new system model. We've learned. We've updated how we think the world works. We've got a new system model.

And we're going to apply that to subsequent decisions. That's the adaptive phase. And really what makes adaptive management adaptive is if you take that learning and apply it to subsequent decisions. That's the adaptation, the adaptation of subsequent decisions as a result of the new information that you've acquired.

Those are the two concepts. That's really what makes recurrent decisions different is they're dynamic and they have the potential to be adaptive.

Now, the interesting question that you get into with regard to this is that, is that learning likely to generate a new understanding and better decisions that are worth the costs of acquiring that learning? And the costs of acquiring that learning are twofold.

One is the monitoring might cost you money. You've got to put the monitoring in place to generate the learning. But the other thing is you might actually take some actions early on that are meant to accelerate learning. And it's an interesting question about whether that accelerated learning is worth, whether the long-term benefits of it are worth, the short term costs, the foregone opportunities of taking that accelerated learning.

I have a question for you to think about. I want you in your small groups to have some discussion and think about examples of recurrent decisions that you are familiar with, that operate in your workplace, perhaps. And I want you to think about how the elements of structured decision making need to be thought of for recurrent decisions.

How do you need to think about objectives perhaps differently for recurrent decisions than for static or single time decisions? Are the actions different? How are the models different? When you're doing the

optimization, how does that need to be different? And when you're doing the monitoring, what's the role of monitoring?

So I want you to take 10 or 15 minutes and think through these ideas, start to put these ideas together that we've talked about in all the modules leading up to this, ideas about objectives, actions, models, optimization, finding the right answer, and now this concept of monitoring, learning, value of information.

How do you need to think about those in recurrent decisions? What opportunities do they provide? And how might you need to think about them differently in recurrent situations rather than static situations? So take some time to discuss that with your group.

OK, I posed quite a challenging question to you actually before you stopped the video. And that was to think about how the elements of structured decision making need to be thought of differently in the case of recurrent decisions where you've got the opportunity for adaptive management.

I'm not actually going to follow up on that discussion in a lot of detail. It opens up a whole new course really. And at the National Conservation Training Center in West Virginia, we teach a course, a five-day course, on adaptive management.

There are also a number of other courses, one that's offered by distance learning quite often, on adaptive management. So there's other training that you can do to get into the nuts and bolts of adaptive management.

I'm going to touch on some of the concepts in the next few minutes but really not get into it deeply. I mean, in some ways, adaptive management is a special case. It is a special case of structured decision making, structured decision making for recurrent decisions.

So all the concepts that we talked about in all the nine modules in this course apply in adaptive management setting as well. You still need to frame the problem in the same way, think about objectives, actions, models, all that kind of stuff. But what's new is the monitoring and the opportunity for adaptation. So let's talk a little bit about that.

First I want to talk about monitoring in the context of management and what we use monitoring for. There's a good article in *Trends in Ecology & Evolution* by Jim Nichols and Ken Williams. And they talk

about the purposes of monitoring in the context of applied natural resource management.

And by applied natural resource management, what we really mean is that we're making decisions. So in the case of a decision context, what do we use monitoring for? And Jim and Ken argue that there's three purposes of monitoring in those applied settings.

Those purposes are, one, to assess the state of the system. Why would we want to assess the state of the system? Well, we might want to assess the state of the system because it might be that the decisions that we're making are state dependent.

That is, if it's a wet year, we'll make one management decision. If it's a dry year, we'll make a different management decision. And depending on those cases, which it is, well, we need to monitor. We need to know whether it's a wet or dry year to know what the best thing to do is.

So monitoring the state of the system in order to make a state dependent decision because there's some decision thresholds that we have and we need to know where we are with regard to those thresholds and depending on where we are we'll make different a decision. So one of the purposes of monitoring is to assess the state of the system.

Second purpose is to determine if our objectives are being met. We're being asked, as conservation managers, as natural resource managers, more and more to be accountable to the public and stakeholders with regard to the goals that we're trying to achieve.

And so one of the ways of being accountable is to take a performance-based approach to things and really monitor the objectives that we're trying to achieve, just as an accounting measure, just to be able to say, did we or did we not achieve these objectives. So that's another purpose for monitoring is to determine if our objectives are being met.

The third purpose for monitoring and the thing that's emphasized the most in adaptive management settings is that we would like to resolve uncertainty, uncertainty that has a high value of information so, therefore, will change subsequent decisions that we make. We pursue that because we think that information is going to allow us to make better decisions in the future. So for that, to reduce that uncertainty, we need monitoring to provide us a feedback in the learning.

But the point I want to make and the point than Jim and Ken make in their article is that the

development of a monitoring system in an applied setting should be tailored to these needs and driven by the decision context. You can't develop a monitoring system for an applied setting without understanding the decision context.

You have to do the structured decision making first. You have to know what the decision is. You have to know what the consequences of uncertainty are. And then you can design the monitoring that's best going to serve these purposes.

So let me give you a little example. This refers back to the example that Sarah presented in module A at the beginning of this course. So recall that you're the manager of a Ponderosa Pine forest. And you're concerned about a number of things. You had objectives. You've got actions. You've got uncertainty.

Well, one of the purposes of monitoring is evaluation just to see if you have met your management objectives. Well, if your management objectives had to do with having an open canopy pine stand with understory vegetation cover of pinegrass, elk sedge, and exotics, then you would need to have a monitoring system that measured whether the canopy closure rate, the understory vegetation cover, the amount of pinegrass, elk sedge, exotics. So the monitoring system would need to have those elements in it if those are your objectives and you want to evaluate achievement of those objectives

Second, suppose you have a management trigger. Suppose you analyze this decision. And it says, well, our management prescription calls for thinning the pine stand when the basal area is greater than 85 square feet per acre. So in any given year, if you're trying to know whether you should implement thinning, you need to know whether the basal area is greater than 85 square feet per acre.

So you need to have a monitoring system that goes out every year and measures the basal area. And the purpose of that monitoring system is to provide the management trigger. OK, that's fine. You understand that management trigger from the decision analysis. And that allows you to design the appropriate monitoring system.

The third purpose for monitoring in the Williams Nichols construct is learning. And so in this particular example that we're using, if you have uncertainty about the effects of mechanical thinning as opposed to prescribed understory fire on vegetation composition and that uncertainty impedes your ability to make a decision, then you might want to adaptively learn about those things.

So you might implement mechanical thinning sometimes and understory fire sometimes. And then you would do the appropriate monitoring that's going to allow you to see the differential effects on the vegetation composition as a result of those different actions.

In other words, in this case, the monitoring is targeted at those measures that will allow you to discern the models that you've got that represent your uncertainty to reduce that uncertainty. So the monitoring is tied to the uncertainty that plagues your management.

In all cases, the monitoring derives its purpose and derives its design from the decision context. And I think that's an important point. Because I think we're happy to go out and design monitoring systems in the absence of a full understanding of the decision context.

And I think what happens there, the cost of that, is that it's an inefficient use of resources. We may be designing monitoring systems that are frankly unnecessary and too costly on the one hand. That's one pitfall.

The other possibility is we might not be designing the right monitoring system so we're not able to achieve the objectives we have for it. So some thought really needs to go into designing a good monitoring system in this applied context of making decisions.

So let's put these concepts altogether and talk about adaptive management itself. So adaptive management is the idea that we're seeking to optimize management decisions in the face of uncertainty using learning at one stage to influence decisions at subsequent stages while considering the acquisition of information in the optimization.

The way I phrased this is actually what's known as active adaptive management. And active adaptive management is adaptive management that's smart about how the decisions we make now might or might not accelerate learning and whether that accelerated learning is worth our while in the long run.

So what does this look like? Well, I got a diagram here that sort of maps out-- everybody's got their own diagram of adaptive management. Here's mine. Start at the node that says Decision right at the T to T plus 1 junction. It's on the top of the diagram on the left side.

So the idea is that we make a decision at some point in time. That decision integrates the actions that we have, the objectives that we have, and our current understanding of how the system works. So

there's a structured decision embedded there.

We make a decision. And we put into place a monitoring system. So we make that decision. We implement it. Then we go out and see what the outcome was. So we do some survey. We gather some other data. That's the monitoring stage of things.

Now we take that monitoring data. We confront our system models with that observed data. Our system models make predictions. The observed data shows us what, in fact, happened. That tension allows us to learn and update our system models.

Maybe we revise the model weights. Maybe we're capturing uncertainty in some different versions of the model and there's weights on those different models and that updating represents our uncertainty. There's other ways to capture learning.

But the idea is that you've got predictions from system models. You've got observations. And the confrontation of those two things produces learning and allows you to update your understanding of how the system works.

Now what do we do? Well, we make new predictions. We look at those predictions relative to our objective function and perhaps our utility functions. We evaluate those for different alternative actions. And we make a new decision and so on.

And we keep doing that time and time again. And that allows us to, as we go, update our understanding of how the system works and adapt to that new information as we make subsequent decisions.

That's the idea of adaptive management then. And again, as I said, there's a lot more to say on the topic. And in fact, there's something interesting about adaptive management I think that's important to note here.

I have a colleague, friend, that I've done some work with. And he works on a very large and very well-known adaptive management problem that's been in place for about 15 years and has generated a lot of learning.

And he said to me recently, he said, Mike, we've been doing adaptive management for 15 years. But we've never been doing the structured decision making. So that's an interesting statement. What is meant by that?

Well, in this particular case, they were focusing on the learning phase of things. But they hadn't really thought about how that learning results in subsequent decisions. So they were focused on the reduction of uncertainty.

But they hadn't folded that into the full decision framework that said, if we reduce the uncertainty, here's how we're going to make subsequent decisions. Here's how our decisions are going to change as a result of that new information.

And that bit, the integration of the learning with the structured decision making, with the decision context, that's where I think you get the full-fledged adaptive management and the adaptive management that can be seen then as a special case of structured decision making.

And so that's really what I want to challenge people to do. The truth is that we throw around the term adaptive management an awful lot. I mean take the *Journal of Wildlife Management* or *Conservation Biology* and do a search on the term adaptive management and see how many hundreds of papers you get.

Virtually everybody says, ah, well, we need to do adaptive management. But look really hard and see in those cases, where people are saying that there's adaptive management going on or arguing for adaptive management, whether they really have this concept of a recurrent decision that's being made that's plagued by uncertainty for which the value of information is high enough that it's worth putting in a monitoring system that's going to allow you to reduce uncertainty over time because that will enhance your decisions in the future and you can even articulate how those decisions will change as a result of that learning.

There's not a lot of articulations of adaptive management that have all those elements folded together. And so I think that's a challenge really to all of us is to do a better job of seeking full-fledged articulations of adaptive management that are really nested deeply in understanding the decision context in which that uncertainty operates.

So I have one more topic to talk about. And that is an issue of problem framing as tied to adaptive management. And so really what we're doing is I'm looping back to the very beginning of this course where we talked about the first proact step and the idea of problem framing.

So as we said in module B, recognizing the core elements of the decision and how they fit together, framing the problem, that really is the hardest part of a decision analysis. And often we don't do it well. We don't really see all the aspects of the decision, perhaps, even until we've implemented it.

Well, now recurrent decisions offer a very interesting opportunity. Because it isn't just that we do a decision analysis. We come up with a preferred alternative, implement it, and then, uh, you know, that's the way it is. It's implemented. It's done. We can't revisit it.

In those cases, if we didn't frame the problem right, then we probably didn't make the best decision. And we only find out about that afterwards when we can't do anything about it.

When we're talking about problem framing for adaptive management problems, do we frame them any differently than we do a run of the mill static problem, a problem that's not recurrent? Do we do problem framing any differently for adaptive management problems than for the broader class of structured decision making problems?

And I think in some ways no. But there are some subtleties. Again, I think the first question is to ask what the decision is. Who's the decision maker? What is the decision? What's the scope and scale, et cetera?

Now, if in that initial problem framing you say, well, this is a recurrent decision, then that tells you that, OK, you might be in the adaptive management world, right? You identify the elements of the decision like you do in any kind of decision problem, the objectives, the actions, the model. And you ask what impedes the decision.

In fact, some of this, I do this always for any kind of problem. I ask what the decision is. I start to develop the elements of the decision, the proact sequence. And I start to say, what impedes this decision?

Sometimes the impediment to the decision is that we don't know the objectives well or we don't know what the possible actions are. We can't make good model predictions.

Sometimes the impediment to the decision is the uncertainty that's inherent in those model predictions, that we don't know what the consequences of our actions are going to be. And that uncertainly makes it difficult for us to choose the right action. That impediment, uncertainty, that's the motivation for adaptive

management.

So if you have a decision that's recurrent and you identify that at the problem framing stage and you also have a decision that's impeded by uncertainty, then you're in the realm of adaptive management. And what you're going to do then in the problem framing and the initial development is identify what that uncertainty is and analyze those different models and look at the value of information and think about how monitoring could be designed to help you reduce that uncertainty and then how you can use that in subsequent decisions.

So the framing of adaptive management problems really starts out the same way as the framing of the broader class of structured decision making problems. But if it's recurrent and if it's impeded by uncertainty, then that leads you into the realm of adaptive management.

When we implement adaptive management, what do we do? Well, we start with a prototype structure. We perform some initial analysis. Maybe we revise that prototype in the course of developing our decision analysis.

But then in an adaptive setting, we implement that decision. We monitor the results. We gain experience implementing that analysis that we did.

But then after some amount of time, that experience we have with working with the system gives us some new insights, not just insights about which model was doing a good job predicting things, but insights at a broader scale about whether we were even pursuing the right objectives in the first place, whether there's a whole different set of actions to consider that we hadn't considered initially, whether there are some new aspects of the system and way we're modeling it that we hadn't considered before that we should consider. And that might really change how we think about the problem.

And sometimes it's really actually difficult to understand the core issues of our problem, too. You've actually implemented a decision. But in the recurrent case that's appropriate in adaptive management, once we've implemented a decision, we're not done. We have the opportunity not only to update our learning about the system but to update perhaps how we even frame the problem.

And this is a concept called double-loop learning. This is a diagram that comes out of the Department of Interior Adaptive Management Guide that was published in 2007. And it captures an idea that's written about in the literature by a lot of people, this idea of double-loop learning.

The idea of adaptive management, one of the ideas of adaptive management, is this iterative phase that, once we've implemented a decision, we can monitor. We can assess the outcome of that decision. We can update our models, do some learning.

And we have what some people call the inner loop of adaptation in the iterative phase here. We can update our decisions year after year based on the learning that accrues.

But that iterative phase also provides us with the experience implementing this decision and may give insights about how we frame the decision in the first place, the set-up phase that really was all that structured decision making kind of stuff, who are the stakeholders, who is the decision maker, what are the objectives, what are the alternatives, what models, what monitoring.

All that set-up stuff that came from our initial decision analysis, maybe we learned some stuff about that that causes us to reflect about how we set things up in the first place. Let me give you an example.

One of one of the preeminent examples of adaptive management in the United States is the management of mallard harvest. This is called adaptive harvest management. It was implemented initially in 1995.

And it's this formal prescriptive system of adaptive management for setting harvest regulations. The Fish and Wildlife Service uses this each year to set the harvest regulations for mid-continent mallards.

And there's uncertainty. It's a recurrent decision. There's uncertainty about how the duck population responds to harvest management. And that uncertainty, there's a monitoring system. That uncertainty's resolved over time. And there's all this optimization that's done, the whole thing.

So this has been implemented in 1995. So at this point, it's been going on for 15, 20, more than 20 years, right, no, not quite 20 years, 15 years, anyway, 15, 16 years. At this point, I mean it's still going. It's still intact. We're still really in the inner loop. We've done some technical updates to the modeling but nothing really major.

But at this point, the water foul management community is involved in some really sort of broad scale discussions about really what are the objectives, not only of water fowl harvest management, but water fowl habitat management under the North American Water Fowl Management Plan. So there's some

very broad scale decisions or discussions at least about what are our fundamental objectives for these management endeavors.

And furthermore, there's a lot of discussion about, what is the uncertainty that really plagues our ability to manage these populations? And there's another question about, have we even framed how we set up these regulations the right way?

Should we be focusing on setting regulations for mallards and then kind of patching together all the other species? Or should we set up regulations that are based on the status of a composite set of duck species? So should we frame the problem quite differently?

So we're in this phase right now where we're rethinking how we frame the decision. I don't know when this might result in changes to how we set regulations in this country. But we're at this sort of double-loop phase of really questioning how we set up the problem.

That's powerful. It's good. And it's one of the opportunities in adaptive management that I think we want to avail ourselves of.

In an interesting note here, I mentioned earlier that there's different schools of thought about adaptive management. People define it differently. There's an interesting article by Jamie McFadden and Drew Tyre and others. And that reference is now flashing on your screen.

And this article actually identifies two schools of adaptive management, two flavors of adaptive management, if you will, the decision theoretics school and the resilience experimentalist school of adaptive management. And there is a common framework that both of these schools of adaptive management have. And you would recognize it. It's a lot of what's painted here in the diagram that I showed and also in this double-loop learning kind of idea.

But what's different about these two schools of thought in some ways is the decision theoretics school emphasizes this structured decision making approach to adaptive management, to the structured decision making at the set-up phase, and to this idea of updating in the iterative phase and adaptation in the iterative phase, focusing on that inner loop of learning and the reduction of uncertainty. That's the decision theoretic school of adaptive management kind of in a nutshell.

The resilience experimentalist school focus much more on the outer loop of this double-loop learning, is

much more interested in using the implementation, the recurrent implementation, of a decision and, perhaps, even experimental management actions to gain information about what we care about, what actions are possible, how we should even set up these models to gain information really about how the problem is framed, who's involved in the problem, what governance structures are needed to implement this decision. And so it's really focused on the outer loop.

And so both of these are really legitimate approaches. And there's things that we need to bring out at different phases of adaptive management. And one will apply more than another in different contexts when we think about things.

So when you stroll through the adaptive management literature and you see these very different approaches, I think it's useful perhaps to have this double-loop picture in mind and understand that these different schools of thought with regard to adaptive management stem from thinking that the concern and where you need to focus your effort is at different phases of the problem. But I think there is a common underlying framework by which we can understand adaptive management.

So it's been my pleasure and Sarah's pleasure to offer this course online through this format. And we thank you very much for your interest and attention.

And we hope that this course has been valuable to you in thinking about how the tools of decision analysis can help you make better decisions as managers of natural resources and how you can move from the intuitive decision making that's based on our inherent abilities as human decision makers to these approaches that are more formal in nature but really in some ways are just a formal embodiment of what we do intuitively anyway but which serve our purposes perhaps in the public eye quite well.

Because they lead to decisions that are transparent, that are explicit, that we're able to include others in the decision and show them where their input is valuable and what we need them for, and what we hope are replicable decisions that are not so much subject to the intuitions and perhaps whims of individual decision makers or individual groups but that, in their transparency and their explicitness, lead to some replicability and which we hope is increasing the quality of decisions and as well the respect that decision making agencies and decision makers earn from the public.

So with that good luck. And feel free to contact us if there's any way that we can help you better understand these concepts. Thanks very much for taking this course.